

Bargaining and Market Behavior

Essays in Experimental Economics

VERNON L. SMITH

University of Arizona



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
10 Stamford Road, Oakleigh, Melbourne 3166, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain

<http://www.cup.cam.ac.uk>
<http://www.cup.org>

© Vernon L. Smith 2000

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2000

Printed in the United States of America

Typeface Times Roman 10/12 pt. *System* QuarkXPress [BTs]

A catalog record for this book is available from the British Library.

Library of Congress Cataloging in Publication Data

Smith, Vernon L.

Bargaining and market behavior : essays in experimental economics / Vernon L. Smith.
p. cm.

Includes bibliographical references.

ISBN 0-521-58450-7

1. Economics – Methodology. 2. Economics – Simulation methods. 3. Economics
– Psychological aspects. I. Title.

HB131.S59 2000

330'.01 – dc21

99-15975

CIP

ISBN 0 521 58450 7 hardback

Contents

<i>Preface</i>	xi
Part I: Economics and Psychology	
INTRODUCTION	3
1. Rational Choice: The Contrast Between Economics and Psychology <i>By Vernon L. Smith</i>	7
2. Experimental Tests of the Endowment Effect <i>By Robert Franciosi, Praveen Kujal, Roland Michelitsch, Vernon L. Smith, and Gang Deng</i>	25
3. Monetary Rewards and Decision Cost in Experimental Economics <i>By Vernon L. Smith and James M. Walker</i>	41
4. Fairness: Effect on Temporary and Equilibrium Prices in Posted-Offer Markets <i>By Robert Franciosi, Praveen Kujal, Roland Michelitsch, Vernon L. Smith, and Gang Deng</i>	61
Part II: Bargaining Theory, Behavior, and Evolutionary Psychology	
INTRODUCTION	79
5. Preferences, Property Rights, and Anonymity in Bargaining Games <i>By Elizabeth Hoffman, Kevin A. McCabe, Keith Shachat, and Vernon L. Smith</i>	90
6. Social Distance and Other-Regarding Behavior in Dictator Games <i>By Elizabeth Hoffman, Kevin A. McCabe, and Vernon L. Smith</i>	127
7. On Expectations and the Monetary Stakes in Ultimatum Games <i>By Elizabeth Hoffman, Kevin A. McCabe, and Vernon L. Smith</i>	139
8. Game Theory and Reciprocity in Some Extensive Form Experimental Games <i>By Kevin A. McCabe, Stephen J. Rassenti, and Vernon L. Smith</i>	152

9. Behavioral Foundations of Reciprocity: Experimental Economics and Psychology *By Elizabeth Hoffman, Kevin A. McCabe, and Vernon L. Smith* 177

Part III: Institutions and Markets

INTRODUCTION	203
10. Reflections on Some Experimental Market Mechanisms for Classical Environments <i>By Vernon L. Smith</i>	206
11. Experimental Methods in the Political Economy of Exchange <i>By Vernon L. Smith</i>	240
12. Individual Rationality, Market Rationality, and Value Estimation <i>By Peter Knez, Vernon L. Smith, and Arlington W. Williams</i>	257
13. Market Contestability in the Presence of Sunk (Entry) Costs <i>By Don Coursey, R. Mark Isaac, Margaret Luke, and Vernon L. Smith</i>	265
14. The Boundaries of Competitive Price Theory: Convergence, Expectations, and Transaction Costs <i>By Vernon L. Smith and Arlington W. Williams</i>	286
15. Off-Floor Trading, Disintegration, and the Bid-Ask Spread in Experimental Markets <i>By Joseph Campbell, Shawn La Master, Vernon L. Smith, and Mark Van Boening</i>	320
16. Bertrand-Edgeworth Competition in Experimental Markets <i>By Jamie Brown Kruse, Stephen Rassenti, Stanley S. Reynolds, and Vernon L. Smith</i>	348
17. An Experimental Examination of the Walrasian Tâtonnement Mechanism <i>By Corinne Bronfman, Kevin A. McCabe, David P. Porter, Stephen Rassenti, and Vernon L. Smith</i>	381

Part IV: Stock Markets and Bubbles in the Laboratory

INTRODUCTION	409
18. Stock Market Bubbles in the Laboratory <i>By David P. Porter and Vernon L. Smith</i>	414
<i>References</i>	434
<i>Index</i>	451

Rational Choice: The Contrast Between Economics and Psychology

Vernon L. Smith

Rational Choice (Hogarth and Reder, 1987) is about economics and psychology, or, as noted by Zeckhauser (1987, pp. 251–4), the rationalist versus behavioralist views of economics. One would have hoped that in this book, given the potential of psychologists and economists to learn from each other, the record would have shown more tangible evidence of this learning.

This chapter discusses the themes in this debate, a debate that is never quite joined: The psychologist's provocative claims are neither answered nor echoed by the economists. My comments will arise from the perspective of experimental economics, which reflects equally the rational and behavioral intellectual traditions. Generally, I want to address the reference to "a growing body of evidence – mainly of an experimental nature – that has documented systematic departures from the dictates of rational economic behavior" (Hogarth and Reder, 1987, p. vii). This suggests a contest between economic theory and the falsifying evidence from psychology. But there is a third view, that of experimental economics, which documents a growing body of evidence that is consistent with the implications of rational models, although there are many important exceptions. In the latter, often the data can be comprehended by modifying the original models. The result is to deepen the concept of rationality and simultaneously increase consistency between the observations and the models; better normative models more accurately predict the experimental results. Psychologists almost uniformly report results contrary to rational theory, which leads them to conclude that the "normative and descriptive analyses of choice should be viewed as separate enterprises" (Tversky and Kahneman, 1987, p. 91).

I. Rationality as Conscious Cognition

"My first empirical proposition is that there is a complete lack of evidence that, in actual human choice situations of any complexity, these [rational] computations can be, or are in fact, performed . . . but we cannot, of course, rule out the

possibility that the unconscious is a better decision maker than the conscious” (Simon, 1955, p. 104).

Throughout much of *Rational Choice*, one sees the frictional tension between psychology and economics. But from the perspective of experimental economics, I believe that the basic problem stems not from the numerous areas of claimed disagreement expressed in *Rational Choice*, but from two unstated premises on which there is implicit agreement between psychology and mainstream theory: (1) rationality in the economy emanates and derives from the rationality of individual decision makers in the economy, and (2) individual rationality is a cognitively intensive, calculating process of maximization in the self-interest. A third shared tenant, which is a correlate of points 1 and 2, is that (3) an acceptable and fundamental way to test economic theory is to test directly the economic rationality of individuals isolated from interactive *experience* in social and economic institutions.¹ Economists do not usually challenge this tenet. They are merely skeptical of the way psychologists implement it: by asking subjects how they would choose among stated hypothetical alternatives. It is reasonable to conjecture from this that the methodology would be acceptable if the decision maker had a “stake” in the decision, in which case the issue could in principle be resolved empirically. But according to point 1, nothing is added or addable by the conjunction of individuals in, and with, markets that cannot be captured by giving the subject a verbal description of the particular market decision-making context. Market rationality is then the direct result of individual choice rationality in that described context.² But experimental economics

¹ Arrow, recognizing point 1 as an implicit assumption in traditional theory, is concerned with correcting this view: “I want to stress that rationality is not a property of the individual alone. . . . It gathers not only its force but also its very meaning from the social context in which it is embedded” (1987, p. 201). But Arrow’s point is about theory: His main theme is that the power of theory derives from the conjunction of rational individuals with the concepts of “equilibrium, competition, and completeness of markets” (p. 203). For example, theory assumes complete information and common knowledge as part of the rationality of individuals, making rationality a social phenomenon.

² Of course, psychologists are interested in studying cognitive processes in decision-making situations that appear to be remote from market processes. But such decisions may still have a social context, such as hospital and medical committees in the case of physician decisions. The study of isolated cognitive processes is of interest in its own right but also needs to be studied explicitly in other social contexts. It is desirable to know whether the strong effect of framing (survival versus mortality probabilities) on physicians’ stated preferences is related to their actual decision to use one therapy rather than another. Presumably, “best practice” therapies evolve in a social context not from isolated individuals thinking about alternatives in terms of probabilities. Experiments that would attempt to capture these social processes would be analogous to the experimental economist’s program of studying market decision making in particular institutional contexts.

research suggests that different results obtain when subjects' choices are interactively governed by an institution. Although the rediscovery of institutions in economic theory began about 1960 with the contributions by Shubik (1959), Coase (1960), Hurwicz (1960), and Vickrey (1961), the new thinking, hypothesizing that institutions matter, is still not well integrated with both theory and laboratory evidence.

That individual rationality is a consciously cognitive phenomena is fundamental in the rhetoric of microeconomic and game theory. The theorist, if called on, says that the model assumes complete information on payoffs (utilities) and more. "The common knowledge assumption underlies all of game theory and much of economic theory. Whatever be the model under discussion . . . the model itself must be assumed common knowledge; otherwise the model is insufficiently specified, and the analysis incoherent" (Aumann, 1987, p. 473). Without such common knowledge people would fail to reason their way to the solution arrived at cognitively by the theorist. This is echoed by Arrow when he notes that a "monopolist, even . . . where there is just one in the entire economy, has to understand all these [general equilibrium] repercussions . . . has to have a full general equilibrium model of the economy" (Arrow, 1987, p. 207). Indeed, it has been hard for either the theorist or the psychologist to imagine optimal market outcomes being achieved by other than conscious cognition; it can't occur by "magic," so to speak.³ The reason is that neither has traditionally modeled markets as a learning process, capable of converging to a rational equilibrium outcome. A noteworthy exception is to be found in Lucas (1987), in which some examples are used to motivate the hypothesis that myopic agents with adaptive expectations converge to steady states, which sometimes correspond to a Muthian rational expectations equilibrium.⁴

What has emerged from 30 years of experimental research is that the preceding premises 1–3 are false. Plott (1987) summarizes many examples. In these experiments (also Smith, 1962), all information on the

³ Simon (1955) is open to the possibility that unconscious decisions may be better than the conscious. But Simon (1987, p. 39) says that "in situations that are complex and in which information is very incomplete (i.e., virtually all real-world situations), the behavioral theories deny that there is any magic for producing behavior even approximating an objective maximization of profits or utilities." Yet there are a great many, very complex, experimental markets, with very incomplete information, that converge to outcomes that precisely approximate those derived from maximizing objectives. We badly need the kind of cooperation between economics and psychology that would help us to better understand how, in Simon's (1987, pp. 26–8) well-known terminology, the procedural rationality of the individual allows substantively rational outcomes to be achieved over time in these markets.

⁴ Other important contributions to the study of the market process are provided by Blume and Easley (1982) and Bray (1982); see also Kalai and Lerher (1993).

economic environment (values) is private; far from having perfect or common information, subjects know only their own “circumstances.” All trading is carried out by an institution such as the decentralized “open outcry” rules of the continuous double auction in which every agent is both a price maker who announces bids to buy (offers to sell) and a price taker who accepts a standing offer to sell (bid to buy). What these and many hundreds of other experiments have shown is that (1) prices and allocations converge quickly to the neighborhood of the predicted rational expectations competitive equilibrium, and (2) these results generalize to a wide variety of posted-price, sealed-bid, and other institutions of exchange, although convergence rates tend to vary and can be influenced by extreme parameter conditions.

Postexperiment discussion with the subjects in the earliest experiments made it plain that (1) subjects are not aware that they are achieving maximum profits collectively and individually, in equilibrium, and, in fact, deny this when asked; and (2) before seeing the results, subjects describe the market situation as confused and disorderly (“How can you get anything out of these experiments?”). When asked what strategies they used, they are unable to convey insight to the experimenter: “I tried to buy low (sell high)” or “I waited until near the end to squeeze the other side.” These and other bidding, auctioning, and price-posting experiments show the predictive power of noncooperative equilibrium concepts (competitive or Nash) without any requirement that knowledge be complete and common. In these cases, economic theory works, predictively, under weaker conditions than expected, and no support is provided for the interpretation that the equilibrating process is consciously cognitive. The verbal behavior of subjects strongly contradicts what their actual behavior achieves.

The fact that private-information experimental markets converge more quickly and reliably to certain rational predictions than complete-information markets do directly contradicts the conclusion of Tversky and Kahneman (1987, p. 88): “Perhaps the major finding of the present article is that the axioms of rational choice are generally satisfied in transparent situations and often violated in nontransparent ones.”⁵ This is correct in their context, but in experimental markets rational theory often performs best in the “nontransparent” (low information) environ-

⁵ According to one of my referees, “Isn’t it odd that one would find this quote [by Tversky and Kahneman] in a book in which Plott demonstrated the operation of a near continuum of markets (the signaling example)? Somehow the psychologists miss the point of examples even when the examples are placed directly in front of them. As I reflect on these papers I do not recall any psychological explanation of any of the papers that have used experimental economics techniques. It seems to me that the psychologists have not done their homework.”

ment and worst in the “transparent” (high information) environment. The leap is so great when one goes from data on responses to individual choice problems to observed behavior in experimental markets that conclusions of this sort are reversed! This underscores the criticism by Coleman (1987) of the implicit premise of the conference that the greatest gains for theory will come from a more sophisticated model of action. “It is deficiencies in the apparatus for moving from the level of the individual actor to the behavior of the system that hold the greatest promise of gain” (p. 184). I think this is the most important implication of experimental economic research. What is imperfectly understood is the precise manner in which institutions serve as social tools that reinforce, even induce, individual rationality. Such economic concepts as noncooperative equilibrium and incentive compatibility are helpful, but they are inexorably static and do not come to grips with the interactive process between agents and institutions. One misses all of this in research limited to the individual expressing an opinion about described situations or alternatives.

It is natural to expect that the unconscious can be a good decision maker only when complexity is absent. The single-market experiments discussed earlier are simple in the sense that there is but one isolated market characterized by stationary supply and demand, but the observed results still follow in some institutions when demand is constantly shifted privately without public announcement of any kind (McCabe et al., 1993). Furthermore, there are many examples showing that in much more complicated multiple-market experiments, convergence to competitive equilibria is observed.⁶

II. Verbal Behavior: Unreliable and Not Worth Studying?

The preceding discussion might lead some to infer, incorrectly, that nothing worth knowing can be learned by studying verbal behavior. Verbal behavior, when studied with the skills of the psychologist, tells one a lot about how people think about choice problems. Their choices

⁶ See Smith (1986, p. 169) and Williams et al. (1986) for examples with three commodities and two markets; see Plott (1988) for an example with 19 connected markets. Another type of complexity occurs in experimental asset markets in which the asset dividend is not only uncertain but also dependent on a sample of likelihood information. It is well known that psychologists find judgment biases that contradict the Bayesian updating of subjective probabilities from sample information. One important study finds that “in eight experiments with inexperienced subjects, prices tend toward the Bayesian predictions, but there is some evidence of exact representativeness bias in prices and allocations. However, the degree of bias is small, and it is even smaller in experiments with experienced subjects. All other non-Bayesian theories can be rejected” (Camerer 1987, p. 995).

deviate from the predictions of rational choice theory, and, as noted earlier and in the studies summarized later, verbal behavior deviates from actual behavior in market experiments. The nature of these discrepancies is a legitimate object of investigation. But to be most informative for economics, such studies should be conducted in the context of experimental (or field) markets, where one observes both verbal intentions and actual choices. The paper by Kahneman et al. (1987) provides an example of what I mean. People, when asked, state that it is fairer to allocate surplus football tickets (above season subscriptions) by lottery or queue than by auctioning to the highest bidders. But what would be the effect on behavior and attitudes toward fairness if the auctioning of tickets makes it possible to lower the price of season tickets or build an addition to the stadium? Economic reality consists of both the unseen and the seen – both the indirect and direct consequences of decision. What economics brings to the analytical table is the broader perspective that choice decisions have repercussions, which actual decision makers experience but may not perceive. These effects are omitted from opinion surveys, but they may have an impact on attitudes as well as behavior. Opinion polls tend to reflect the psychologist's direct perception of reality and not the economist's view that "decisions have consequences." This is why secondary effects and equilibrium concepts are an integral part of the designs used by experimental economists.⁷

Another example of the insights to be gained by embedding fairness questions into a broader decision context has been provided by Hoffman and Spitzer (1985b). They study Coase (1960) bargaining between pairs of subjects, one of whom has been given a position of advantage assigned by a coin flip. All subjects bargain to an efficient solution, but none of the privileged bargainers receives the larger individually rational share afforded by the property right. To see what might be driving these results, Hoffman and Spitzer replicated their experiments with the following change: Instead of by a coin flip, the position of advantage is awarded to the winner of a game of Nim, and both subjects are told that the winner of this pregame will have "earned" the right to be the controller in the bargaining game. Now two-thirds of the controllers negotiate individually rational bargains. This finding suggests that assigning rights at random induces a very strong unintentional fairness ethic that largely disappears when the rights are earned. Because rights in the economy do not arise by random assignment but are acquired through some eco-

⁷ I am not criticizing the idea put forward by Kahneman et al. (1987) that norms of fairness may be necessary to explain certain data. The point is that there are naïve and sophisticated notions of "fairness"; also, in experimental markets, self-interested behavior tends to overcome any fairness norms.

conomic or political process, the interpretation of fairness opinion research is ambiguous without a deeper examination of the circumstances.⁸ Thus Kahneman et al. find that most people consider it fair for a firm to lower price if input costs decline; but if costs decline because of an invention by the firm, then it is fair to retain the resulting profit. (They earned it?) But do these opinions have anything to do with observed behavior in a market? There is abundant evidence that prices fall, over time, in experimental markets when the induced supply schedule is increased. Kahneman et al. could answer this question by first soliciting people's opinions, as discussed earlier, and then doing market experiments under two different announced explanations for the cost reduction: an input price decline and an innovation. I would bet on no significant difference, in equilibrium, although initially I think it is plausible that market prices will be higher in the second treatment. The point is that these are researchable questions, and we can ask how verbal behavior relates to motivated behavior in experiments. Fairness is likely to be most important in understanding the regulated firm: how regulation came to be and how it operates politically.

Some research that combines verbal and actual market behavior has been reported by experimental economists. I will summarize three studies: (1) verbal, or choice, versus actual marketlike behavior in the context of preference reversal phenomena; (2) the relation between opinion polls and a presidential stock market; and (3) laboratory stock markets in which verbal behavior does not contradict and helps to explain market behavior.

A. Preference Reversals and Markets

The standard preference reversal problem is easily described. The subject chooses gamble A over B (or B over A) but then states that his or her minimum willingness-to-accept price for A is less (more) than the minimum willingness-to-accept price for B. Such choice-revealed preference reversals vary from 24 to 68% of subject samples, with monetary rewards making it in a subject's interest not to preference reverse (Chu

⁸ There is also a problem with all the bargaining literature in which there is a first-mover advantage (see Kahneman et al., 1987, for references). This literature reports overwhelming evidence that a bargainer's utility functions include fairness norms, but as far as I have been able to determine, in all such cases, the first-mover advantage is awarded at random. The usual interpretation of the Hoffman-Spitzer result is that it shows how fairness is sensitive to framing. But what it shows is much more significant: that equal-split bargaining results may be due, generically, to an important treatment thought to be benign, namely, the standard use of random devices to allocate subjects to initial conditions.

and Chu, 1990, p. 906). These results have been replicated many times, with a great variety of gambles going back to the path-breaking work of Lichtenstein and Slovic (1971). Chu and Chu (1990) construct a marketlike con game for subjects in which each is asked to state a preference between two gambles, with the understanding that if later the subject were holding a gamble that was not a preferred one (inferior or no preference), then a “trader” (one of the experimenters) could exercise the right to exchange the other gamble for the one held. Because each subject’s exchange price was also solicited, the experimenter was in a position to return to the earlier choice and arbitrage the subject’s holding. Up to three arbitrage transactions were performed against preference-reversing subjects in a sequence of decisions and trades. The results dramatically reduced the frequency of preference reversals: “three transactions were all that was needed to wipe out preference reversals completely” (Chu and Chu, 1990, p. 909). Arbitrage makes rational choice transparent.

B. The Iowa Presidential Stock Market

Comparative data on verbal versus market behavior are provided by the presidential stock market conducted by experimentalists at the University of Iowa (Forsythe et al., 1991). In return for an initial minimum expenditure of \$35, subjects received ten \$2.50 portfolios consisting of ten shares in each candidate (at the time, Bush, Michael Dukakis, Jesse Jackson, and the rest of field) plus a \$10 deposit credit from which further net purchases in the market were deducted. All funds invested were paid to shareholders in the form of a single dividend payment accruing at 9:00 A.M., November 9, 1988, when the market closed. The dividend paid to each share held on that date was \$2.50 times the proportion of the popular vote cast for each candidate.⁹

In the Iowa experiments, Bush’s lead as predicted by the market (Bush price – Dukakis price)/\$2.50, rose to 8 points by October 20 and varied from November 2–9 within the range 6.8–8. The final results showed Bush winning by a margin of 7.8%. During the same period, Bush’s margin in the opinion polls varied from 4 to 14%. The market was both more accurate and less volatile than the polls’ measurement of

⁹ Because uncertainty in this experiment was not of the usual balls-in-urns variety, it relates directly to the critique by Einhorn and Hogarth (1987), who “believe that it is time to move beyond the tidy experiments and axiomatizations built on the explicit lottery” (p. 64). This view is discussed by Zeckhauser (1987, pp. 257–8). Of course, in all market experiments, subjects face uncertainty about the behavior of others, but subject values either are certain or are generated by lotteries. Forsythe et al. (1991) is an innovative exception.

verbal behavior based on standard representative sampling techniques. The subject sample was not, however, representative; it consisted of those who bought presidential portfolios (restricted to the University of Iowa community). Furthermore, opinion polls were not “news” to the market. The frequent publication of poll results had little if any discernible effect on market prices. Hence, the well-known inaccuracy of political opinion polls was appropriately discounted by the market. Five surveys of candidate preferences and beliefs among the stock market participants revealed judgment biases of the kind found by psychologists and political scientists in the study of opinion. But Forsythe et al. were unable to identify significant patterns of trading behavior due to presidential preferences. For example, male–female differences in preferences for Bush, as expressed in polls, are large but were not reflected in differences between male and female holdings of Bush; the greatest difference was 4.2%, which was not statistically significant.

C. Laboratory Stock Markets: Verbal Behavior Can Complement Actual Behavior

I shall close this section with a discussion of laboratory stock market experiments, in which verbal behavior does not contradict, and complements, actual behavior by helping to provide an “explanation” of behavior in terms of subjects’ perceptions of their experience. Dozens of laboratory stock market experiments have documented that price bubbles arise naturally from the home-grown capital gains expectations of subjects in environments in which probabilistic dividend value is common information in each trading period of the horizon (Smith et al., 1988; Schwarz and Ang, 1989). These bubbles are particularly pronounced with first-time subjects, show a reduced amplitude and trading volume when subjects return for a second market experiment, and essentially disappear in a third session in which prices deviate much less from fundamental dividend value and volume is very light. Shares ultimately trade near their rational expectations value, but it takes three 15-period trading sessions in a stable environment to yield this equilibrium. These results do not inspire confidence in rational expectations theory in asset markets where the determinants of share value are not common information.

The verbal behavior of subjects following a bubble experiment is revealing. Some are puzzled by the failure of shares to trade at fundamental dividend value and with the “panic buying” they observe. Many report amazement at the speed with which a market crash can occur; they had expected to sell out ahead of the others when the crash came.

After the market turns, some are hesitant to sell because they cannot bring themselves to take the capital loss or because they hope for a recovery. Many report a reluctance to sell before the crash because they were “too greedy.” Somehow, the volatile behavior of the market was due to the other traders. Although they have no causal explanation of their experience (prices rise “without cause”) and their consensus forecasts never predict the crashes, their comments are consistent with the market observations, with a self-reinforcing expectations view of the boom, and with the tendency of the market crash to dividend value to take two or three periods to occur.

Why is verbal behavior helpful to the researcher in this environment and so obviously misleading in the supply and demand experiments discussed previously? The primary difference, I think, resides in the complete common information characteristics of the first environment and the private information state of the second. Subjects understand the structure of the stock market experiments, but they have divergent expectations and are uncertain about the behavior of others (until they come to have common expectations). Consequently, they are all drawn into myopic behavior. They are aware of this and reveal it in their verbal responses. In the supply and demand experiments, they are not aware of a structure to understand beyond their private circumstances and therefore are unable to relate to what happens in the market, which is unfathomable. If this interpretation is correct, it suggests that the solicitation of verbal responses is more likely to be of interest and to provide insights to the researcher in games of complete rather than private information. This is unfortunate because the world is more likely to conform to the private- than to the complete-information environment (Shubik, 1959, p. 171).

III. On the “Little” Evidence That Monetary Rewards Matter

Monetary incentives are commonly absent in the research of psychologists. This has made their work vulnerable to the criticism that the results are not meaningful; rational choice, it is argued, is a theory that assumes high stakes. This is a specious argument in that it is not offered as a formal part of the theory and is invoked informally only upon finding that the theory does poorly in the face of the experimentalist's data. But it is a legitimate question, and so long as psychologists avoid systematic (not casual) comparisons of the effect of payoffs on decision, they avoid joining this issue. Instead, what we see is the claim that there is “little” evidence of improvement when monetary rewards are introduced (Tversky and Kahneman, 1987, p. 90; Thaler, 1987, p. 96). It is unclear

how seriously one should take this claim when it is converted immediately into a straw man with statements of the form, “the evidence that high stakes do not always improve decisions” (Tversky and Kahneman, 1987, p. 90) or “the assertion that systematic mistakes will always disappear if the stakes are large enough” (Thaler, 1987, p. 96). I know almost no one who would use the adverb *always* in these contexts. In support of the claim – Grether and Plott (1979) is the popular citation, but strangely missing, because it is said that there *is* a little evidence – are any of the many citations that could have been offered showing that monetary rewards matter. Plott (1987, p. 120) is more eclectic, and credible, in citing evidence on both sides because his work includes a case in which money didn’t matter (individual choice) and a case where it did (committee decisions).

A considerable number of experimental studies have assessed the effect on decisions of rewards versus no rewards, and/or varying the level of rewards. Seventeen such studies showing increased support for rational models with increased rewards, and new data on auction behavior when rewards are varied from 0 to 20 times the level customarily used, are reported in Smith and Walker (1993a, b). These studies include a wide variety of institutions and environments: Bernoulli choice decision (see following discussion), bilateral bargaining, Cournot oligopoly, Bertrand oligopoly, casino betting, and double-auction trading. In almost all cases, one effect of increased payoffs is to reduce the standard error of the observations around the predicted or estimated optimal decision. These results are consistent with decision cost models that postulate a trade-off between the benefits of better decisions and the subjective cost of making them.

IV. Is It Rational To Be “Rational”?

“To predict how economic man will behave we need to know not only that he is rational, but also how he perceives the world – what alternatives he sees and what consequences he attaches to them.” (Simon, 1956, p. 271)

Siegel (1959, 1961) appears to be the first psychologist to respond to Simon’s suggestion that rational behavior needs to be examined from the perspective of the individual’s perception and experience of the decision situation. Simon’s thinking was ultimately transformed into concepts of “satisficing” and “bounded rationality,” which were secondary interpretations of the more fundamental characteristic of humankind quoted earlier. Anyone familiar with the rational choice paradigm who thinks along implementational lines will realize that the subjective cost of exploring options and figuring out what to do must be part of the

problem of rational choice as experienced by the decision maker. Decision cost is the cost of concentration, attention, information acquisition, thinking, monitoring, checking, deciding, and acting – all the things you do to realize a decision. When the benefits are small, the decision cost may not be worth it, or the decision cost incurred may be, rationally, correspondingly small so that the pain fits the pleasure, marginally speaking, although the typical subject will not consciously think about it in such terms. Alternatively, in familiar or repetitive situations the decision process may be so automatic, programmed, and instinctive that decision cost is nothing. The decision cost model does not imply that introducing rewards will always improve normative performance, but merely that there is a bounded trade-off between the cost of decision and its value (Smith, 1976a). For example, increasing the monetary stakes may have little effect in a signal detection experiment where the subject's hearing capacity is already strained to its physiological limit. It is therefore somewhat surprising that some signal detection experiments report improved performance with increased reward (Swets and Sewell, 1963, pp. 123–4; Calfee, 1970, pp. 898–9); improvement is often not statistically significant because improvement is small relative to variability for each individual. But the number of such individuals dominates the sample.

Siegel (1961) reexamined the Bernoulli trials experiment. In this task, the subject is importuned to “do your best” to predict, on each trial, which of two events will occur. One event was programmed to occur with probability p , so that its complement would occur with probability $1 - p$. The subjects made repeated choices over a great many trials (up to 1,000; see, e.g., Edwards, 1961). The standard result was “probability matching”; that is, the mean prediction was approximately p for the frequent event. The standard conclusion was that people were not “rational” because one “should” choose the frequent event 100% of the time if one wants to maximize the expected number of correct predictions. Siegel wondered, since there was no monetary or other explicit reinforcement for getting it “right,” whether perhaps the results were not the exception that proved the maximizing rule. Utility theory does not predict that people will maximize something when it is not in their interest to do so. He also pondered what the beleaguered subjects saw in this long string of repeat choices and conjectured that they must be incredibly bored. So he hypothesized in the form of a model that individuals diversified their predictions to relieve boredom. In the absence of monetary rewards, the model was consistent with probability matching; with monetary rewards, the model predicted a shift toward choosing the frequent event more often, and the greater the incentive the greater the shift. Siegel ran experiments with different reward levels, with conditions that afforded sub-

jects the opportunity to relieve boredom without the necessity of diversifying choice, with adult subjects, and with children. In all cases the model qualitatively predicted the observed shifts in the data (Siegel, 1959, 1961; Siegel et al., 1964).

Others (Swensson, 1965; Tversky and Edwards, 1966) also report Bernoulli choice experiments comparing monetary rewards with “do your best” instructions, showing that mean responses shifted upward on the frequent event when rewards were used. In the light of the preceding discussion, the Tversky-Edwards case is noteworthy for the reason that the authors discounted the significance of the increased support for maximization under the reward condition. They conclude that, although the results were in the direction predicted by the normative model, they were “far indeed” from the predictions of that model (p. 682). In closing, they conjecture that a formal model of the results might introduce a “cost associated with making a decision” (p. 683), as indeed it might, and did, 7 years earlier! It seems that today psychologists who “argue that the deviations of actual behavior from the normative model are too widespread to be ignored, too systematic to be dismissed as random error, and too fundamental to be accommodated by relaxing the normative system” (Tversky and Kahneman, 1987, p. 68) are able to do so by ignoring contrary interpretations and evidence over extended periods of time.

If the experimenter is open to the possibility, subjects teach us something about rationality that is not part of standard theory. A good example concerns the matter of incentive compatibility in the (two-sided) sealed-bid-offer auction where each buyer and seller has multiple differing values for units in the aggregate market demand and supply schedules. It is well known that such auctions are not incentive-compatible: If all agents submit fully revealing bids and offers except the buyer with the marginal valuation, then that buyer can increase personal profit on intramarginal units by underrevealing the marginal unit and setting a lower price. Because of this incentive incompatibility, theory has supposed that outcomes cannot be efficient. In privacy experiments, no one knows who has the marginal units, but in repeat play in stationary environments, subjects soon discover that one or more of their units are near the realized “crossing” prices and that any person determining price can manipulate it to the advantage of his or her side of the market. What subjects learn in repeat play is to submit almost all their bids to buy and offers to sell at prices very near or equal to the estimated auction price, where the latter tends to converge to the competitive equilibrium. This causes the reported supply and demand schedules to be elastic but to cross in the neighborhood of the true crossing, with many tied bids. This unconscious group strategy has the property that each side is protected

from manipulation by the other side; any seller whose offer is equal to the auction price and who (thinking that she is determining the price) offers at a higher price next period simply gets replaced by a tied alternative seller with no change in the auction price. This is best described as superrational: a better strategy than has been deduced from rational theory and one that clearly violates the behavioral implications of the standard impossibility theorems.

V. Sunk Costs and Opportunity Costs

According to the usual interpretation of economic theory, the optimizing firm or individual should ignore sunk costs in weighing gains and losses at the margin and should treat out-of-pocket costs as equivalent to opportunity costs. These are strong predictions, based on extremely simple models of maximization, and there is no shortage of verbal measures of behavior showing that these predictions are not confirmed (Thaler, 1980; Kahneman et al., 1987). At the conscious level, the disutility of loss may stimulate unusual effort to recover sunk costs, whereas opportunity costs may require a sophistication in decision analysis that is not quickly attained without formal training. But, as we have already seen, it is common for experimental markets to approximate optimal results in the absence of participant understanding of the private or social conditions that define their market situation. Plott (1987, pp. 122–5) presents experiments in which buyers and suppliers representing the basic demand and supply environment are separated, with intermediary traders purchasing units from the suppliers at one location and then reselling them separately to the buyers at another. These results showed no evidence of market failure due to the sunk cost fallacy. The same is true of a much larger number of intertemporal competitive equilibrium experiments going back many years (Miller et al., 1977; Forsythe et al., 1982; Williams and Smith, 1984). In these experiments, traders who have acquired inventory to carry over for resale at a different place or time tend to take their lumps when they have paid more than they can recover in the resale market. This doesn't mean that there are never any individual cases of irrationality; it means only that such behavior is not important enough to vitiate the static competitive model in these experiments.

Experimental markets for two interdependent commodities (Smith, 1986, p. 169; Williams et al., 1986) do not support the implication that the opportunity cost fallacy upsets the empirical validity of the competitive model in this more complicated setting. In these experiments, the demand curve for each commodity is conditional upon the price (oppor-

tunity cost) of the other commodity. Yet these markets converge to the competitive equilibrium. The convergence is somewhat slower than in single markets, and many subjects do not reach the exact maximizing commodity two-tuple. But those who buy too much of one or the other of the two commodities tend to be offset by those who buy too little. In this case, individual errors are not rare, but they tend to be offsetting so that the market predictions of the model are very good. Economists sometimes suggest that irrationalities lead to random error and can be ignored, but psychologists are right in asserting that this is an empirical question requiring evidence.

The opportunity cost fallacy in the psychology literature often arises in the context of fairness and framing: "For example, most respondents believe that it is unfair for a store to mark up the jars of peanut butter in its stock when wholesale prices rise, apparently because they associate the cost to the individual jar" (Kahneman et al., 1987, p. 113). It is traditional retail practice to relabel shelf stock when the wholesale price rises. Does this mean that retail store managers must "understand" the opportunity cost principle? No, it means merely that this institutional tradition encodes the social learning that results when, historically, managers have tried policies other than what the economist calls the opportunity cost principle. A simple real scenario illustrates this process. Suppose the policy is not to change the price of shelf stock either when wholesale prices rise or when they fall or both. This policy doesn't work for price declines. Either the manager must wait until all the existing shelf stock is sold, in which case she has empty space unutilized, or she must have merchandise marked with two different prices. Price-conscious customers will note this and take the cheaper stock so that the manager will have trouble selling the older stock. It becomes outdated, and customers complain that the stock is not fresh. A customer who did not notice the different prices discovers in the checkout line that somebody else paid less and complains to the manager that this is not "fair." The policy also does not work for rapid price increases. I observed this in California in 1974. The price of sugar rose very rapidly (enough so that cafeterias, rationally, stopped putting sugar packets on the tables; one had to ask for sugar at the checkout counter, where it could be monitored; and managers did not have to think, "monitoring cost is now lower than the cost of pilfered sugar!"). The supermarkets had succumbed to consumer activist group pressure, and whenever wholesale prices rose, they followed the policy of not marking up the price of shelf stock. This created the temporary situation in which the retail sugar price was below the wholesale price. Bakers, candy makers, and other wholesale buyers discovered this, and cleaned out the retail shelves. This situation is hardly

fair, but it is a predictable consequence when prices are not permitted to equilibrate. The point is that pricing rules evolve in response to experience, not logical analysis, and policies that are disequilibrating (causing the manager trouble) are altered; this process of alteration continues until we have a tradition. The current manager does not know about opportunity cost or even why the policy is what it is; only that she learned it from the last manager. She is an instrument of the “law” of one price in a market, which is not a law at all, but a tendency that gets encoded in markup relabeling policies. Because psychologists are not, by training, fully informed on these matters, they ask oversimplified questions. The appropriate question is whether it is “fair” to depart from the traditional policy of remarking shelf stock in response to price changes when such departures cause the kinds of problems enumerated here. One does not ask the appropriate question because no one understands all the ramifications. One cannot learn much of anything about economic process by limiting oneself to these kinds of verbal observations. One can, of course, understand something about people’s perceptions, but a model of the system can’t be built around such observations. What they do help us to see is how simple-minded and unworkable interventionist policies get on the political agenda.

VI. Criticism of Experimental Economics

Psychologists since Siegel have not attempted to apply their perspective and questions to market experiments of the kind studied by experimental economists, although there are numerous examples of research by experimental economists who have examined questions studied by psychologists. Mutual understanding is therefore limited, as can be inferred from the two criticisms of experimental economics contained in Hogarth and Reder’s introduction to their book. The first notes that schemes designed to provide incentive compatibility add complexity, and “it is not clear that subjects understand the full implications of such reward structures” (1987, p. 12). They fail here to understand why experimental economists use incentive-compatible rules. We want to see whether such rules make a difference in the observations when compared with procedures that do not use them (as in comparing uniform price and discriminative price institutions). There is no presumption in these comparisons that people “understand the full implications” of such rules. We do not expect such understanding to occur, certainly not in the sense that theorists can be said to understand incentive compatibility. But “understanding” need not have and often does not have anything to do with effectiveness, as we have observed in thousands of experiments. The question comes back

to the conditions under which people achieve socioeconomic outcomes that are not part of their conscious intention; our data tell us that they do. Psychologists do not study this question because psychologists are avowedly, rightly, and appropriately concerned to study cognitive processes. Their experiments are designed for and limited to this worthy objective; our experiments have rather different objectives, and we would not claim that our subjects understand all or any of their situation.

The second criticism of experimental economics is that the introduction of constraints determined by the limit-price supply and demand constraints “places such severe restrictions on subjects’ actions that psychologists may wonder whether the underlying economic theory is in fact being tested” (Hogarth and Reder, 1987, p. 12). This criticism reveals ignorance of the large literature in which limit-price constraints are not used – auctions, multiple markets, asset markets, and supply and demand markets with uncertain redemption values (Plott and Agha, 1983).

Experimental economics can benefit greatly from the criticisms of psychologists, but in order for this to occur, their knowledge and understanding of the literature and its motivation will have to move beyond the superficial levels of familiarity exhibited in *Rational Choice*.

VII. Concluding Discussion

Why is it that human subjects in the laboratory frequently violate the canons of rational choice when tested as isolated individuals but, in the social context of exchange institutions, serve up decisions that are consistent (as if by magic) with predictive models based on individual rationality? Experimental economists have no good answers to this question, although adaptive learning models such as those of Lucas (1987) are suggestive. We need the help of psychologists, undeflected by battles with straw men. It seems evident that an important part of the answer resides in the properties of exchange institutions and how privately informed, but globally poorly informed, decision making is mediated by institutions. Although institutional rules, since Vickrey (1961), have attracted increasing interest among theorists, that development has been slow. In any case, we tend to analyze institutions as given in history not to address the question of why institutions have the structure and rules we observe. I want to suggest that perhaps the structures we observe have survived because of their merit in coaxing Pareto-efficient behavior out of agents who do not know what that means.

Language learning in children occurs in a social context (Brown, 1973). Without contact with people, children do not learn to speak. If

they have such contact, they learn to speak in the total absence of formal instruction. But the same can be said of decision making: I could substitute “make market decisions” for “speak” in the last two sentences and they would apply to what we have learned in the laboratory about adults. On the basis of cognition alone, without the language of the market and ongoing social interaction with other agents, rational decision is frustratingly illusive.